



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/752,874	12/29/2000	Lily P. Looi	2207/9869	1458

7590 03/20/2007
KENYON & KENYON
Suite 600
333 W. San Carlos, Street
San Jose, CA 95110-2711

EXAMINER

PATEL, NIMESH G

ART UNIT	PAPER NUMBER
----------	--------------

2111

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
2 MONTHS	03/20/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

MAILED

MAR 20 2007

Technology Center 2100

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 09/752,874
Filing Date: December 29, 2000
Appellant(s): LOOI ET AL.

Paul E. Steiner
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed November 20, 2006 appealing from the Office action mailed May 17, 2006.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

5,987,538	Tavallaei et al.	11-1999
5,944,809	Olarig et al.	8-1999
6,119,191	Neal et al.	9-2000
6,189,065	Arndt et al.	02-2001

Intel, Multiprocessor Specification, Version 1.4, May 1997, Pages 2-2 to 3-16

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

U35 U.S.C. 102(b) Rejection

Claims 1-3, 9-12, 16-21 and 25-26 are rejected under 35 U.S.C. 102(b) as being anticipated by Tavallaei et al.(US Patent 5,987,538), hereinafter referred to as Tavallaei.

Regarding claim 1, Tavallaei discloses an interrupt delivery system that has a first pair of scaleable node controllers(Figure 2, Component 14; Component 14 is scaleable since the number of Components 14 can change), wherein each node controller supports at least 1 microprocessor(Figure 2, Components 12). Tavallaei further discloses a first scalability port switch(Figure 2, Component 26; Component 26 is scalable since the number of interrupts handled can increase) coupled to each of said scaleable node controllers, wherein said first scalability port switch is to receive an interrupt request, determine an ID, i.e. address, of one of said scaleable node controllers from said interrupt request and transmit said interrupt request to said one of said scaleable node controllers(Column 7, Lines 6-9, 54-56; Column 8, Lines 43-46). Tavallaei shows all of the elements recited in claim 1 and therefore, claim 1 is rejected.

Regarding claim 2, Tavallaei discloses an interrupt delivery system, further comprising a peripheral component interconnect device(Figure 1, Component 34). Tavallaei shows all of the elements recited in claim 2 and therefore, claim 2 is rejected.

Regarding claim 3, Tavallaei discloses an interrupt delivery system, further comprising a peripheral component interconnect bus coupled between the peripheral component interconnect device and the first scalability port switch, wherein said peripheral component interconnect bus is able to support a plurality of additional peripheral component interconnect device(Column 6, Lines 47-53; Component 26 and 28 are integrated and therefore Component 26 is also

Art Unit: 2111

connected to the PCI bus 32). Tavallaei shows all of the elements recited in claim 3 and therefore, claim 3 is rejected.

Regarding claim 9, Tavallaei discloses a method for delivering an interrupt request in a multi-node computer system, comprising: receiving an interrupt request at a scalability port switch(Figure 2, Component 26; Component 26 is scalable since the number of interrupts handled can increase); determining an ID, i.e. address of a scaleable node controller(Figure 2, Component 14; Component 14 is scaleable since the number of Components 14 can change and the system can still function well) to receive said interrupt request; and transmitting said interrupt request to said scaleable node controller(Column 7, Lines 6-9, 54-56; Column 8, Lines 43-46). Tavallaei shows all of the elements recited in claim 9 and therefore, claim 9 is rejected.

Regarding claim 10, Tavallaei discloses a method for delivering an interrupt request in a multi-node computer system, further comprising determining a processor to receive the interrupt request(Column 7, Lines 6-9). Tavallaei shows all of the elements recited in claim 10 and therefore, claim 10 is rejected.

Regarding claim 11, Tavallaei discloses a method for delivering an interrupt request in a multi-node computer system, further comprising comparing a priority of the interrupt request with a priority of the processor(Column 7, Line 41-44). Tavallaei shows all of the elements recited in claim 11 and therefore, claim 11 is rejected.

Regarding claim 12, Tavallaei discloses a method for delivering an interrupt request in a multi-node computer system,, further comprising interrupting the processor(Column 7, Lines 6-9). Tavallaei shows all of the elements recited in claim 12 and therefore, claim 12 is rejected.

Regarding claim 16, Tavallaei discloses an interrupt request that is generated by a PCI device(Column 4, Lines 62-63). Tavallaei shows all of the elements recited in claim 16 and therefore, claim 16 is rejected.

Art Unit: 2111

Regarding claim 17, Tavallaei discloses a method, wherein the interrupt request is generated by a processor(Column 4, Line 50). Tavallaei shows all of the elements recited in claim 17 and therefore, claim 17 is rejected.

Regarding claim 18, Tavallaei discloses a set of instructions residing in a storage medium, said set of instructions capable of being executed by a processor for delivering an interrupt request in a multi-node computer system, comprising: receiving an interrupt request at a scalability port switch(Figure 2, Component 26; Component 26 is scalable since the number of interrupts handled can increase); determining an ID, i.e. address of scaleable node controller(Figure 2, Component 14; Component 14 is scaleable since the number of Components 14 can change and the system can still function well) to receive said interrupt request; and transmitting said interrupt request to said scaleable node controller(Column 7, Lines 6-9, 54-56; Column 8, Lines 43-46). Tavallaei shows all of the elements recited in claim 18 and therefore, claim 18 is rejected.

Regarding claim 19, Tavallaei discloses a method for delivering an interrupt request in a multi-node computer system, further comprising determining a processor to receive the interrupt request(Column 7, Lines 6-9). Tavallaei shows all of the elements recited in claim 19 and therefore, claim 19 is rejected.

Regarding claim 20, Tavallaei discloses a method for delivering an interrupt request in a multi-node computer system, further comprising comparing a priority of the interrupt request with a priority of the processor(Column 7, Line 41-44). Tavallaei shows all of the elements recited in claim 20 and therefore, claim 20 is rejected.

Regarding claim 21, Tavallaei discloses a method for delivering an interrupt request in a multi-node computer system,, further comprising interrupting the processor(Column 7, Lines 6-9). Tavallaei shows all of the elements recited in claim 21 and therefore, claim 21 is rejected.

Art Unit: 2111

Regarding claim 25, Tavallaei discloses an interrupt request that is generated by a PCI device(Column 4, Lines 62-63). Tavallaei shows all of the elements recited in claim 25 and therefore, claim 25 is rejected.

Regarding claim 26, Tavallaei discloses a method, wherein the interrupt request is generated by a processor(Column 4, Line 50). Tavallaei shows all of the elements recited in claim 26 and therefore, claim 26 is rejected.

Claims 1-3, 9-12, 14-21, and 23-26 are rejected under 35 U.S.C. 102(b) as being anticipated by Olarig et al.(US Patent 5,944,809), hereinafter referred to as Olarig.

Regarding claim 1, Olarig discloses an interrupt delivery system that has a first pair of scaleable node controllers(Figure 4, Component 107, 306; More components of type 107 and 306 can be used and therefore it is scaleable), wherein each node controller supports at least 1 microprocessor(Figure 4, Components 105, 106). Olarig further discloses a first scalability port switch(Figure 4, Component 312; Column 8, Lines 8-12; Since more components of type 312 can be used, it is scalable) coupled to each of said scaleable node controllers, wherein said first scalability port switch is to receive an interrupt request, determine an address of one of said scaleable node controllers from said interrupt request and transmit said interrupt request to said one of said scaleable node controllers(Column 8, Lines 35-40; Column 9, Lines 57-65; LOPICs are bus agents that inherently have an address. By delivering the interrupt request to the destination LOPIC, an address has to be determined to deliver the interrupt request to the proper LOPIC). Olarig shows all of the elements recited in claim 1 and therefore, claim 1 is rejected.

Regarding claim 2, Olarig discloses an interrupt delivery system, further comprising a peripheral component interconnect device(Since Olarig discloses a PCI bus(Figure 4, 113;

Art Unit: 2111

Column 8, Line 1), it is inherent Olarig's system comprises a PCI device). Olarig shows all of the elements recited in claim 2 and therefore, claim 2 is rejected.

Regarding claim 3, Olarig discloses an interrupt delivery system, further comprising a peripheral component interconnect bus(Figure 4, 113) coupled between the peripheral component interconnect device and the first scalability port switch, wherein said peripheral component interconnect bus is able to support a plurality of additional peripheral component interconnect device(Column 8, Lines 16-17; Since there are plurality of I/O devices, and there is a PCI bus, there are plurality of PCI devices). Olarig shows all of the elements recited in claim 3 and therefore, claim 3 is rejected.

Regarding claim 9, Olarig discloses a method for delivering an interrupt request in a multi-node computer system, comprising: receiving an interrupt request at a scalability port switch(Figure 4, Component 312; Column 8, Lines 8-12; Since more components of type 312 can be used, it is scalable); determining an address of a scaleable node controller(Figure 4, Component 107, 306; More components of type 107 and 306 can be used and therefore it is scaleable) to receive said interrupt request; and transmitting said interrupt request to said scaleable node controller(Column 8, Lines 35-40; Column 9, Lines 57-65; LOPICs are bus agents that inherently have an address. By delivering the interrupt request to the destination LOPIC, an address has to be determined to deliver the interrupt request to the proper LOPIC). Olarig shows all of the elements recited in claim 9 and therefore, claim 9 is rejected.

Regarding claim 10, Olarig discloses a method of determining a processor to receive an interrupt request(Column 9, Lines 57-65), Olarig shows all of the elements recited in claim 10 and therefore, claim 10 is rejected.

Art Unit: 2111

Regarding claim 11, Olarig discloses a method of comparing a priority of the processor's task with that of the interrupt request(Column 3, Lines 4-7; Column 10, Lines 8-10). Olarig shows all of the elements recited in claim 11 and therefore, claim 11 is rejected.

Regarding claim 12, Olarig discloses a method for interrupting the processor(Column 10, Lines 10-13). Olarig shows all of the elements recited in claim 12 and therefore, claim 12 is rejected.

Regarding claim 14, Olarig discloses an interrupt request that is a broadcast interrupt(Column 9, Lines 56-67, and Column 10, Lines 1-7; A distributed delivery mode interrupt is a broadcast interrupt). Olarig shows all of the elements recited in claim 14 and therefore, claim 14 is rejected.

Regarding claim 15, Olarig discloses the use of an end of interrupt register to indicate the end of processing of an interrupt(Column 7, Lines 50-55). Olarig shows all of the elements recited in claim 15 and therefore, claim 15 is rejected.

Regarding claim 16, Olarig discloses an interrupt request that is generated by a PCI device(Column 8, Line 17; Since Olarig discloses a PCI bus(Figure 4, 113; Column 8, Line 1), it is inherent Olarig's system comprises a PCI device that will generate an interrupt). Olarig shows all of the elements recited in claim 16 and therefore, claim 16 is rejected.

Regarding claim 17, Olarig discloses a method, wherein the interrupt request is generated by a processor(Column 7, Lines 28-30). Olarig shows all of the elements recited in claim 17 and therefore, claim 17 is rejected.

Regarding claim 18, Olarig discloses a set of instructions residing in a storage medium, said set of instructions capable of being executed by a processor for delivering an interrupt request in a multi-node computer system, comprising: receiving an interrupt request at a scalability port switch(Figure 4, Component 312; Column 8, Lines 8-12; Since more components

Art Unit: 2111

of type 312 can be used, it is scalable); determining an address of a scaleable node controller(Figure 4, Component 107, 306; More components of type 107 and 306 can be used and therefore it is scaleable) to receive said interrupt request; and transmitting said interrupt request to said scaleable node controller(Column 8, Lines 35-40; Column 9, Lines 57-65; LOPICs are bus agents that inherently have an address. By delivering the interrupt request to the destination LOPIC, an address has to be determined to deliver the interrupt request to the proper LOPIC). Olarig shows all of the elements recited in claim 18 and therefore, claim 18 is rejected.

Regarding claim 19, Olarig discloses a method of determining a processor to receive an interrupt request(Column 9, Lines 57-65), Olarig shows all of the elements recited in claim 19 and therefore, claim 19 is rejected.

Regarding claim 20, Olarig discloses a method of comparing a priority of the processor's task with that of the interrupt request(Column 10, Lines 8-10). Olarig shows all of the elements recited in claim 20 and therefore, claim 20 is rejected.

Regarding claim 21, Olarig discloses a method for interrupting the processor(Column 10, Lines 10-13). Olarig shows all of the elements recited in claim 21 and therefore, claim 21 is rejected.

Regarding claim 23, Olarig discloses an interrupt request that is a broadcast interrupt(Column 9, Lines 56-67, and Column 10, Lines 1-7; A distributed delivery mode interrupt is a broadcast interrupt). Olarig shows all of the elements recited in claim 23 and therefore, claim 23 is rejected.

Regarding claim 24, Olarig discloses the use of an end of interrupt register to indicate the end of processing of an interrupt(Column 7, Lines 50-55). Olarig shows all of the elements recited in claim 24 and therefore, claim 24 is rejected.

Art Unit: 2111

Regarding claim 25, Olarig discloses an interrupt request that is generated by a PCI device. generates an interrupt request(Column 8, Line 17; Since Olarig discloses a PCI bus(Figure 4, 113; Column 8, Line 1), it is inherent Olarig's system comprises a PCI device that will generate an interrupt). Olarig shows all of the elements recited in claim 25 and therefore, claim 25 is rejected.

Regarding claim 26, Olarig discloses a method, wherein the interrupt request is generated by a processor(Column 7, Lines 28-30). Olarig shows all of the elements recited in claim 26 and therefore, claim 26 is rejected.

35 U.S.C. 103(a) Rejection

Claims 4-5 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tavallaei, in view of Neal et al.(US Patent 6,119,191), hereinafter referred to as Neal.

Regarding claim 4, Tavallaei discloses of an input/output hub(Figure 1, Component 28; Devices are connected to component 28 and therefore can act as an input/output hub) coupled between the PCI bus and the port switch. Tavallaei does not disclose multiple PCI hubs connected to the input/output hub. However, Neal discloses multiple PCI hubs that are connected to a hub(Figure 5). Therefore, it would have been obvious to one of ordinary skill in the art to combine the teachings of Tavallaei with the teachings of Neal because this would allow more PCI devices to be connected:

Regarding claim 5, Tavallaei discloses a second pair of node controllers coupled to a switch(Figure 1, Component 14).

Art Unit: 2111

Claims 6-7 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tavallaei and Neal as applied to claims 4-5 above, and further in view of Multiprocessor Specification, hereinafter referred to as MP.

Regarding claims 6 and 7, Tavallaei and Neal do not disclose the use of an additional switch connected to the first input/output hub. However, MP discloses the use of multiple switches(Figure 2-2; I/O APIC). Therefore it would be obvious to combine the teachings of Tavallaei and Neal with the teachings of Olarig to have a second port switch connected to the first input/output hub because it would increase interrupt scalability(Page 3-13).

Claims 13 and 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tavallaei, in view of Arndt et al.(US Patent 6,189,065), hereinafter referred to as Arndt.

Regarding claim 13, Tavallaei does not disclose the method of redirecting the interrupt request to a different processor. However, Arndt discloses redirecting an interrupt to a different processor(Column 8, Claim 8). Therefore it would have been obvious to one of ordinary skill in the art to combine the teachings of Arndt and Tavallaei to redirect an interrupt to different processor since this would allow an interrupt to be handled faster than waiting for the original, busy processor to handle the interrupt.

Regarding claim 22, Tavallaei does not disclose the method of redirecting the interrupt request through the scalability port switch to a different processor. However, Arndt discloses redirecting an interrupt to a different processor(Column 8, Claim 8). Therefore it would have been obvious to one of ordinary skill in the art to combine the teachings of Arndt and Tavallaei to redirect an interrupt to different processor since this would allow an interrupt to be handled faster than waiting for the original, busy processor to handle the interrupt.

Art Unit: 2111

Claims 13 and 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Olarig, in view of Arndt.

Regarding claim 13, Olarig does not disclose the method of redirecting the interrupt request to a different processor. However, Arndt discloses redirecting an interrupt to a different processor (Column 8, Claim 8). Therefore it would have been obvious to one of ordinary skill in the art to combine the teachings of Arndt and Olarig to redirect an interrupt to different processor since this would allow an interrupt to be handled faster than waiting for the original, busy processor to handle the interrupt.

Regarding claim 22, Olarig does not disclose the method of redirecting the interrupt request to a different processor. However, Arndt discloses redirecting an interrupt to a different processor (Column 8, Claim 8). Therefore it would have been obvious to one of ordinary skill in the art to combine the teachings of Arndt and Olarig to redirect an interrupt to different processor since this would allow an interrupt to be handled faster than waiting for the original, busy processor to handle the interrupt.

(10) Response to Argument

In response to Appellant's argument that the Examiner misconstrues the claim term "scaleable node controller" and that Tavallaei's local APIC 14 of figure 2 is not identical to a scalable node controller, Examiner respectfully disagrees. The specification filed December 29, 2000 discloses a scaleable node controller connected to a SPS 30 and processors 24a-d (Page 5, Lines 12-14; Figure 2). The specification further discloses that a scaleable node controller sends an interrupt request to a processor (Page 5, Line 30). No further evidence has been submitted that clearly defines a scaleable node controller. With the information provided in the specification, a scaleable node controller functions as a component transmitting an interrupt request to a processor. Tavallaei discloses a local advanced programmable interrupt

Art Unit: 2111

controller(APIC) 14 transmitting an interrupt(INTR) to a processor(Figure 2). Therefore, Appellant's scaleable node controller and Tavallaei's local APIC function equivalently, viewed in light of the specification. Appellant's arguments require more functions than the disclosed scaleable node controller. Thus, Examiner's interpretation of a scaleable node controller is fair and reasonable. Examiner also notes that Appellant has had ample opportunity to amend the claims to include the aspects that Appellant desires to distinguish.

Further, a recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use, then it meets the claim. In a claim drawn to a process of making, the intended use must result in a manipulative difference as compared to the prior art. See *In re Casey*, 152 USPQ 235 (CCPA 1967) and *In re Otto*, 136 USPQ 458, 459 (CCPA 1963).

In response to Appellant's argument that the Examiner misconstrues the claim term "scalability port switch" and that Tavallaei's I/O APIC 26 of figure 2 is not identical to a scalability port switch, Examiner respectfully disagrees. The specification discloses a scalability port switch connected to SNCs 22 and 26 and IOH 32(Page 5, Lines 12-13; Figure 2). The specification further discloses that a scalability port switch sends data to a SNC(Page 5, Lines 22-23). No further evidence has been submitted that clearly defines a scalability port switch. With the information provided in the specification, a scalability port switch functions as a component transmitting data to a SNC. Tavallaei discloses an I/O APIC 26 transmitting data to an APIC 14, i.e. SNC(Figure 2). Therefore, Appellant's scalability port switch and Tavallaei's I/O APIC function equivalently, viewed in light of the specification. Appellant's arguments require more functions than the disclosed scalability port switch. Thus, Examiner's interpretation of a

Art Unit: 2111

scalability port switch is fair and reasonable. Examiner also notes that Appellant has had ample opportunity to amend the claims to include the aspects that Appellant desires to distinguish.

Further, a recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use, then it meets the claim. In a claim drawn to a process of making, the intended use must result in a manipulative difference as compared to the prior art. See *In re Casey*, 152 USPQ 235 (CCPA 1967) and *In re Otto*, 136 USPQ 458, 459 (CCPA 1963).

In response to appellant's arguments of the Examiner's theory of inherency is legally deficient in regards to the rejection based on Tavallaei, Examiner respectfully disagrees since inherency is not relied upon to reject claim 1. Tavallaei discloses a scalability port switch, i.e. I/O APIC, determining an address of one of said scalable node controllers from said interrupt request (Column 7, Lines 7-9, 54-56; Column 8, Lines 43-46, The destination field, i.e. address, is determined and the interrupt data is delivered to an appropriate APIC, i.e. scalable node controller). Examiner also notes that the specification discloses the scalability port switch transmitting the interrupt to the scalable node controller based on node ID (Page 5, Lines 22-24). Examiner cannot find anywhere in the specification determining an actual address of the scalable node controller.

In response to appellant's arguments of Tavallaei failing to teach the limitation of comparing a priority of the interrupt request with a priority of the processor, Examiner respectfully disagrees. Tavallaei discloses priorities associated with interrupts, which processor it is directed to (Column 7, Lines 41-44) and therefore the interrupt priority is compared with that of the processor and the task it is performing. It is inherent an interrupt request can only interrupt the processor if the interrupt request has a higher priority than the processor task being

Art Unit: 2111

performed. This is further supported in the background section of the specification of the current application(Page 2, Lines 13-15).

In response to Appellant's argument that the Examiner misconstrues the claim term "scaleable node controller" and that Olarig's combined circuit of cache 107 and LOPIC 306 in figure 4 is not identical to a scalable node controller, Examiner respectfully disagrees. The specification discloses a scaleable node controller connected to a SPS 30 and processors 24a-d(Page 5, Lines 12-14; Figure 2). The specification further discloses that a scaleable node controller sends an interrupt request to a processor(Page 5, Line 30). No further evidence has been submitted that clearly defines a scaleable node controller. With the information provided in the specification, a scaleable node controller functions as a component transmitting an interrupt request to a processor. Olarig discloses combined circuit of cache 107 and LOPIC 306 transmitting an interrupt(INTR) to a processor(Figure 4). Therefore, Appellant's scaleable node controller and Olarig's combined circuit of cache 107 and LOPIC 306 function equivalently, viewed in light of the specification. Appellant's arguments require more functions than the disclosed scaleable node controller. Thus, Examiner's interpretation of a scaleable node controller is fair and reasonable. Examiner also notes that Appellant has had ample opportunity to amend the claims to include the aspects that Appellant desires to distinguish.

Further, a recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use, then it meets the claim. In a claim drawn to a process of making, the intended use must result in a manipulative difference as compared to the prior art. See *In re Casey*, 152 USPQ 235 (CCPA 1967) and *In re Otto*, 136 USPQ 458, 459 (CCPA 1963).

Art Unit: 2111

In response to Appellant's argument that the Examiner misconstrues the claim term "scalability port switch" and that Olarig's COPIC 312 of figure 4 is not identical to a scalability port switch, Examiner respectfully disagrees. The specification discloses a scalability port switch connected to SNCs 22 and 26 and IOH 32 (Page 5, Lines 12-13; Figure 2). The specification further discloses that a scalability port switch sends data to a SNC (Page 5, Lines 22-23). No further evidence has been submitted that clearly defines a scalability port switch. With the information provided in the specification, a scalability port switch functions as a component transmitting data to a SNC. Olarig discloses a COPIC 312, i.e. scalability port switch, transmitting data to combined circuit of Cache 107 and LOPIC 306, i.e. SNC (Figure 4). Therefore, Appellant's scalability port switch and Olarig's COPIC function equivalently, viewed in light of the specification. Appellant's arguments require more functions than the disclosed scalability port switch. Thus, Examiner's interpretation of a scalability port switch is fair and reasonable. Examiner also notes that Appellant has had ample opportunity to amend the claims to include the aspects that Appellant desires to distinguish.

Further, a recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use, then it meets the claim. In a claim drawn to a process of making, the intended use must result in a manipulative difference as compared to the prior art. See *In re Casey*, 152 USPQ 235 (CCPA 1967) and *In re Otto*, 136 USPQ 458, 459 (CCPA 1963).

In response to appellant's arguments that Examiner's theory of inherency is legally deficient in regards to Olarig failing to teach a scalability port switch to determine an address of one of said scalable node controllers from said interrupt request, Examiner respectfully disagrees. Olarig discloses delivering interrupt data to a LOPIC, i.e. scalable node controller,

Art Unit: 2111

and therefore would be determining the address of the LOPIC(Column 9, Lines 62-64). Appellant points out that LOPICs contain a who-am-I register, which maintains a processor ID. By delivering the interrupt request to the LOPIC, this processor ID, i.e. address, has to be determined to deliver the interrupt request to the proper LOPIC. Examiner also notes that the specification discloses the scalability port switch transmitting the interrupt to the scalable node controller based on node ID(Page 5, Lines 22-24). Examiner cannot find anywhere in the specification determining an actual address of the scalable node controller.

In response to appellant's arguments of Olarig to teach the limitation of comparing a priority of the interrupt request with a priority of the processor, Olarig discloses the comparison of priorities of the processor and the interrupt(Column 3, Lines 4-7; Column 10, Lines 8-10). The specification of the current application discloses comparing priority of the processor as "compares the priority of the IRQ with the priority of the current processor task"(Page 8, Lines 11-12). Therefore Olarig discloses the method of comparing priority of the processor.

In response to Appellant's argument of Tavallaei failing to teach a first input/output hub, Examiner respectfully disagrees. The specification discloses an input/output hub connected to a SPS 30 and P64H, which is connected to (Page 5, Lines 12-13; Figure 2). No further evidence has been submitted that clearly defines an input/output hub. With the information provided in the specification, an input/output hub functions as a component transmitting interrupt requests to a scalability port switch. Tavallaei discloses I/O chipset 28, i.e. input/output hub, that transmits interrupt requests to I/O APIC 26, i.e. scalability port switch(Figure 4). Therefore, Appellant's input/output hub and Tavallaei's I/O chipset 28 function equivalently, viewed in light of the specification. Appellant's arguments require more functions than the disclosed input/output hub. Thus, Examiner's interpretation of a scaleable node controller is fair and reasonable. Examiner also notes that Appellant has had ample opportunity to amend the claims to include the aspects

Art Unit: 2111

that Appellant desires to distinguish. Examiner further notes that Examiner has not suggested replacing I/O chipset 28 with Neal's hubs, as alleged by the Appellant.

Further, a recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use, then it meets the claim. In a claim drawn to a process of making, the intended use must result in a manipulative difference as compared to the prior art. See *In re Casey*, 152 USPQ 235 (CCPA 1967) and *In re Otto*, 136 USPQ 458, 459 (CCPA 1963).

In response to Appellant's traversal of the rejection of claims 6 and 7, Examiner respectfully disagrees. Appellant simply states that Neal and MP fail to make up the deficiencies in Tavallaei and seems to be relying on the references individually. One cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986). As stated in the rejection above, Tavallaei and Neal do not teach the use of an additional switch connected to the first input/output hub. However, MP discloses the use of multiple switches(Figure 2-2; I/O APIC). Therefore it would be obvious to combine the teachings of Tavallaei and Neal with the teachings of Olarig to have a second port switch, i.e. scalability port switch, connected to the first input/output hub because it would increase interrupt scalability(Page 3-13).

In response to Appellant's traversal of the rejection of claims 13 and 22, Examiner respectfully disagrees. Again, Appellant seems to be relying on the references individually. One cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986). As stated in the

Art Unit: 2111


rejection above, neither Tavallaei nor Olarig teach a scalable node controller redirecting the interrupt request through the scalability port switch to a different processor. However, Arndt teaches redirecting an interrupt to another processor(Column 8, claim 8). It would have been obvious to one of ordinary skill in the art to use the teachings of Arndt, in the system of Tavallaei and the system of Olarig to redirect the interrupt request through the scalability port switch to a different processor since this would allow an interrupt to be handled faster than waiting for the original, busy processor to handle the interrupt.

(11) Related Proceeding(s) Appendix

The Examiner is unaware of any related proceedings.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,





Nimesh Patel

Examiner


MARK H. RINEHART
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100

Conferees:


MARK H. RINEHART
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100
REHANA PERVEEN
SUPERVISORY PATENT EXAMINER
3/15/07